

METHOD AND DEVICE FOR OPERATING A SECONDARY OPERATING SYSTEM
IN ADDITION TO A PRIMARY OPERATING SYSTEM

The invention relates to a method and a device for operating a secondary operating system on a computer in addition to a primary operating system.

5

The operation of two operating systems loaded on a computer, more precisely in the random access memory of a computer, not only alternatively, but alternately without restarting the computer is known.

10

Thus, WO 98/09225 discloses an operating system for real time extension for conventional, intrinsically not real timeable, Microsoft Windows systems through special microkernels.

15

DE 44 06 094 C2 also discloses a real time extension of conventional Microsoft Windows systems by means of a complete real time operating system, which can also be run separately, i.e. independently of Windows on a computer. The secondary real time operating system has direct access to a single subset of the processor register and hardware components of the computer.

20

It is also known to emulate under an operating system a virtual computer or machine on which a second operating system can run. The secondary operating system is under the control of the monitor program emulating the virtual computer. The secondary operating system cannot directly access all the processor registers, but instead this only takes place under the control of the monitor program. It is problematical in the prior art that the source code of at least one of the operating systems must be known, because during operation in parts thereof not disclosed normally by the supplier intervention takes place or changes must be made to such parts and in particular the same parts on at least one of the operating systems. It is also disadvantageous if one operating system must run "under" another, i.e. is embedded in the latter.

The problem of the invention is to provide a method and a device by means of which at least two operating systems can be run on a computer without reducing their performance and in particular whilst maintaining real time capacities without intervening in the secondary operating system and at the most with interventions in its board support package. However, when the secondary operating system is active it operates on the central processor unit (CPU) as if it was loaded as the sole operating system and can therefore access the entire processor and its virtual memory area without any restriction.

According to the invention the set problem is solved by a method of the aforementioned type, which is characterized in that a secondary operating system driver (SOS driver) of

the primary operating system is loaded for loading and controlling the secondary operating system.

A device for operating a secondary operating system on a computer in addition to a primary operating system therefore provides a secondary operating system driver (SOS driver) of the primary operating system for loading and controlling the secondary operating system.

10 According to the invention access of the secondary operating system takes place without the aid of the primary operating system.

In the development where both operating systems are implemented in juxtaposed manner on a computer means that both operating systems operate completely independently of one another thereon and in particular the operating systems do not presuppose one another, i.e. the secondary operating system does not presuppose the operable presence of the primary operating system. The system driver is exclusively used for changing operating systems. Thus, the entire security area of the POS, including the system driver, could be overwritten into the SOS without impairing its function, but then a return to the POS is naturally impossible. This more particularly implies that one operating system (particularly the secondary operating system) is not set up on the other, particularly the primary operating system or would only access the same. During the operation of one operating system no code part of the other is executed. In particular, during its operation, the secondary operating system does not access the system driver. Thus, according

to the invention, no one operating system is embedded in the other, i.e. does not permanently presuppose the same. This more particularly applies to the secondary operating system relative to the primary operating system. For polling and running one operating system, particularly the secondary operating system, there is no need for any informations from the other operating system, namely the primary operating system, including the system driver in the first-mentioned operating system (secondary operating system).

Both operating systems and in particular also the secondary operating system run in a completely autonomous manner.

As a result of the inventive provision of a driver of the primary operating system for driving and loading the secondary operating system via its board support package, it is also ensured that there is no need to modify the core or kernel of the secondary operating system for the operation thereof as a secondary operating system alongside a primary operating system.

20

The board support package is the software forming the connection between hardware (i.e. the board) and an operating system (i.e. the support). Operating systems used on several platforms (hardware environment, including processor, memory, etc.) always have a BSP and which is consequently a fixed component of the operating system. Embedded operating systems, such as Windows CE, comprise an operating system kernel and the BSP, through whose modification the operating system can be and must be adapted to a specific hardware platform without having to know the operating system kernel.

30

The board support package of an operating system is regularly disclosed by the operating system supplier in the source code, because it in particular contains the so-called basic hardware services by means of which for the operating system in question the necessary interfaces are provided to the hardware, such as e.g. interrupt controllers and system timers and which have been developed for this purpose for use on different hardware systems, i.e. different CPU platforms, in the same way as manufacturer-specific hardware, which in particular differs from the quasi-standard. Thus, through the solution according to the invention it is possible to use all operating systems as a secondary operating system, which with the aid of a board support package can be configured and adapted without any loss to the performance thereof when used as a secondary operating system and, to the extent that they are real time operating systems, without losing their real time capacity.

20

In a preferred development, the inventive method provides for the fact that on changing the dependence of the operating systems there is a replacement of interrupt tables in the volatile memory. Thus, the device according to the invention is constructed to this end. Thus, as a result of the informations provided by the interrupt table entries of the secondary operating system during an interrupt in the secondary operating system the correct interrupt service routine of the secondary operating system is started, so that from this standpoint the process is the same as if there was no primary system. Thus, in this way it is pos-

30

sible to operate in parallel or side by side two operating systems without knowledge of the source code. Thus, according to the invention, in this way no primary operating system information is stored in the secondary operating system memory area. Neither operating system has any information on the other system.

In a preferred development of the method according to the invention, the secondary operating system driver (SOS driver) loaded with the primary operating system loads the secondary operating system in a memory area of the physical random access memory and preferably the upper area thereof not used by the primary operating system.

According to a preferred development, in the computer processing unit (CPU) are created memory contexts (virtual operating areas) and in particular the SOS driver can set up in the CPU a tunnel context by means of which it is possible to switch a change to the operation of the operating systems. The virtual operating area referred to as a context comprises random blocks of the physical memory. A memory management unit (MMU) manages such contexts in a memory allocation table called MMU table and which is used for writing the context. From the programming standpoint operation takes place in the virtual address area referring to the physical memory through the operation of the MMU.

According to a further development, following the loading of the secondary operating system an entry takes place into the same and more precisely into the board support package,

which in a further development sets up a context for booting the secondary operating system in the processor.

According to a further development, into the tunnel context
5 is inter alia loaded a tunnel memory page contained in the driver and into which the program sequence is branched. Then, by means thereof, program codes of the secondary operating system are loaded into the new memory context and the complete booting process of the secondary operating system is continued.
10

For performing the aforementioned method steps, according to a preferred development of the inventive device, the SOS driver has a SOS loading section and a tunnel area containing the tunnel memory page.
15

In a preferred development of the inventive method, after loading or during any operation of the secondary operating system, there is a change from the latter to the primary
20 operating system either when the secondary operating system is idling (entry thereof into its idle loop) or through a corresponding return instruction in the program sequence of the secondary operating system for a return to the primary operating system.

25

For performing this method step, in the device according to the invention the board support package has a corresponding return section.

30 In another preferred development of the inventive method, there is a change from the primary operating system to the

secondary operating system as a result of a SOS interrupt request intended for the secondary operating system.

For this purpose the SOS driver has an interrupt table section by means of which it generates in the primary operating system an interrupt polling or call table (interrupt table), which inter alia contains a poll of the interrupt servicing routine for polling the secondary operating system.

10

In a preferred development of the inventive method, an interrupt servicing routine in the SOS driver reads the interrupt table of the secondary operating system and the processing of the latter takes place or is continued at the point concerning the interrupt poll or call.

15

If there is an interrupt call not intended for the primary operating system, the SOS system driver intercepts it and transfers it via the board support package to the secondary operating system, so that according to the invention there is a change between the operating system by means of the primary operating system SOS driver and the secondary operating system board support package.

20

More particularly and as far as a processor does not support the direct change of contexts, according to a further development of the invention, the change in activity of the operating systems takes place by means of a tunnel area in the tunnel context set up in addition to the primary operating system and also the secondary operating system context.

25

30

According to a preferred development of the inventive method, on changing from one operating system to the other all the system states (in the random access memory) are
 5 stored, particularly all the CPU registers and preferably in addition all the processor-internal caches are emptied. In a corresponding further development on changing from one operating system to the other, the new system states of the other operating system, particularly CPU register contents
 10 and memory management (MMU) tables are loaded into the processor.

Finally, according to a further development of the invention, the timing for the secondary operating system takes
 15 place through the main hardware timer, i.e. only the secondary operating system has access thereto, whilst the timing for the primary operating system takes place through a timing system driver. This is in particular provided and is advantageous if the secondary operating system is a real
 20 time operating system for controlling an industrial plant or machine, whereas the primary operating system is operated by an operator and makes available to him for operational purposes an ergonomic, graphic surface.

25 Further advantages and features of the invention can be gathered from the claims and the following description of an embodiment of the invention with reference to the attached drawings, wherein show:

30 Fig. 1 The allocation of primary and secondary operating systems to the individual

resources of a computer and the design of an operating system using the example of the secondary operating system.

5 Fig. 2 A diagram concerning the process of loading the secondary operating system and, when the latter is idle, the return to the primary operating system.

10 Fig. 3 A change from the primary operating system to the secondary operating system.

Fig. 4 A representation for the synchronization of the clock of the primary operating system.

15

Fig. 1 shows the allocation of the hardware resources (HR) of a computer to the primary operating system (POS) and to the secondary operating system (SOS) within the scope of the present invention. It also shows the fundamental design of an operating system using as the example the secondary operating system SOS.

20

An operating system firstly has a kernel K as the central core thereof. It also has a board support package (BSP) in which are implemented the basic hardware services (BHS).

25

For the operating system the latter provide the necessary interfaces to the hardware, such as to the interrupt controller, system timer, etc. The BHS make it possible for the operating system to be used on different hardware systems.

30

Normally the BSP is made available in source form

by the manufacturer, so that the hardware manufacturer can adapt the operating system to his hardware.

The operating system has further architecture services (As), which form the interfaces necessary for the operating system to the CPU-specific services, such as exception and interrupt processing and MMU management. The As make it possible for the operating system to operate on different CPU platforms.

10

The operating system also has generic operating system services (OSS), which make available application software AS (AS1, AS2), high-quality services, such as memory management, network services, multitasking services, etc. In an operating system normally only the BSP is available in source form, whereas adaptations of architecture and generic operating system services are impossible. It is also possible to see the physical memories PS1 and PS2 associated with the two operating systems, together with parts of the interrupt controls IC1 and IC2 controlling the interrupts for the POS and SOS.

15
20

As can be gathered from fig. 1, important resources of the computer for the operating system are allocated in the following way:

25

At the time of their activity, both operating systems use the entire virtual memory (VM) and all the CPU registers (CPU-R), such as in particular the standard register, floating point register, control register, etc. Both operating systems initially share the random access memory (RAM),

30

each operating system having part of said memory and the secondary operating system (SOS) is preferably loaded into the upper part of the RAM. They also share the interrupt controller (IRC), clear interrupts being allocated to each
5 operating system.

The secondary operating system has in particular sole access to the system timer (ST), whilst the primary operating system (POS) has sole access to the hard disk (HD). To the
10 particular operating system are allocated drivers T1 (of the primary operating system) or T2 (of the secondary operating system), inter alia for managing the particular allocated additional hardware AHP for the primary operating system and AHS for the secondary operating system.

15

One of the drivers of the primary operating system is the SOS-T driver for loading the secondary operating system.

Fig. 2 illustrates the loading of the secondary operating
20 system (SOS).

It is firstly assumed that the primary operating system (POS) is loaded conventionally by means of a boot loader into the computer memory and that the system driver SOS-T
25 for the secondary operating system (SOS) also has been loaded.

After loading the primary operating system (POS), the SOS-T firstly loads the secondary operating system (SOS) into a
30 separate memory area of the RAM not used by the POS and preferably the upper area of the physical RAM, so that the

secondary operating system driver SOS-T forms the boot loader for the SOS.

Besides the memory context for the primary operating system (POS) set up in the conventional way, for the use of the secondary operating system the system driver firstly sets up a context BC for booting the secondary operating system and a tunnel context TC in the CPU for changing between the contexts PC of the primary operating system and the BC.

10

Entry from the system driver SOS-T into the secondary operating system takes place following the loading process in the board support package of the SOS (steps 1 and 2 in fig. 2). As is indicated by the arrow, the system driver SOS-T is exclusively required for the change from the primary to the secondary operating system. The program processing is then continued in the secondary operating system board support package. The board support package sets up the context BC required for SOS booting. The program sequence is then branched into tunnel area TA of the tunnel context TC (step 3), the boot context BC is loaded (step 4) and the secondary operating system booting process is continued (step 5). The SOS sets up its own context in which branching takes place during each future POS to SOS change.

For this purpose the tunnel area contains the "switching code" between the boot loader or secondary context and the tunnel context. The tunnel area comprises precisely one memory page in which is filed the switching code. In all contexts (tunnel context and boot loader or secondary context) said switching code is located at the same virtual address.

After changing to the secondary operating system (SOS) and during the operation thereof, the system driver SOS-T is no longer required by the secondary operating system and does not access the latter or other parts of the primary operating system (POS) and also does not make use of the same. To this extent the secondary operating system operates completely autonomously.

10 If the secondary operating system (SOS) enters an idle state or loop (IL), there is automatically a return to the primary operating system (steps 6, 7 and 8). For as long as the secondary operating system is active, interrupt calls which appear are exclusively processed by the SOS.

15

Prior to loading the secondary operating system (SOS) and during each change to it, all the registers of the CPU primary operating system are stored and the CPU-internal caches are emptied.

20

On loading the secondary operating system and on changing to it, all the secondary operating system CPU registers and the MMU tables thereof are loaded.

25 On changing from the secondary operating system (SOS) to the primary operating system (POS), there is a saving of the system states of the secondary operating system and the loading of system states for the primary operating system in a corresponding way.

30

If the central processing unit (CPU) supports a direct change of contexts, such as is e.g. the case with the Intel-X86 architecture, via task-safe segments, there is also no need for the tunnel context, so that the change between operating systems can take place directly via the primary context (PC) and secondary context (SC).

A change of activity of the primary operating system (POS) for polling and for activity of the secondary operating system (SOS) takes place exclusively as a result of an interrupt intended for the latter. If this occurs the system driver (SD) and the BSP of the secondary operating system are branched into the memory context (MC) of the secondary operating system (SOS) and run there the corresponding allocated interrupt service routine.

During activity of the primary operating system for a change to the secondary operating system, the steps represented in fig. 3 are performed:

20

If the primary operating system is active and there is an interrupt, the CPU as a result of the interrupt request branches directly into the system driver (step A) and the interrupt is intercepted by the system driver. The system driver (SD) saves all the processor registers and changes to the tunnel context (TC) (fig. 2). The system driver branches directly into generic interrupt processing of the secondary operating system BSP. Then, directly by means of or through access to the SOS interrupt table it is established where the interrupt service routine (ISR) of the

30

secondary operating system is located and to where processing is to branch (step B).

From the system driver via the tunnelling function activating the secondary operating system context (SC) (step C),
5 the generic interrupt processing branches directly into the interrupt processing function (ISR) in the secondary operating system (SOS) (step D) and not via any specific code of the SOS, which is not required and is not present according
10 to the invention.

The secondary operating system (SOS) is then run until all the processes deliver their computing time and are consequently branched into the idle loop of the SOS.

15

When the SOS is active and an interrupt occurs for the same, without further interventions, there is automatically a polling of the interrupt service routine (ISR). The processor jumps to this, because during the change from the
20 primary operating system in conventional manner to the secondary operating system the interrupt table in the processor is changed, i.e. during said change the interrupt table of the secondary operating system (SOS) is set. Thus, fig. 3 only shows step D. It is consequently also possible to
25 cancel the primary operating system (POS) and its system driver, without there being any limitations to the secondary operating system (SOS).

From the idle loop and via the tunnel area there is a change
30 to the tunnel context (TC) (steps 6 and 7 in fig. 2). The tunnel function then returns to the system driver of

the primary operating system (POS) (step 8), from where it again activates the primary context (PC) and continues the operation of the POS.

- 5 As an alternative to the use of the idle loop (IL) of the secondary operating system (SOS) as the entry point for changing to the primary operating system, it is possible to initiate said change using a regular process in the SOS. For this purpose the SOS has a function which can be polled
10 by processes. Polling then branches back again to the primary operating system (steps 6 to 8) until the next interrupt for the SOS occurs.

As stated hereinbefore, the main system timer (MST) is controlled by the secondary operating system (SOS). For this
15 purpose the primary operating system is so modified, e.g. by patches, that all accesses to the main system timer (MST) are intercepted by the system driver (SD). The latter stores informations as to the clock rate for operating
20 the primary operating system (POS). The clock rate in the secondary operating system must be higher than that of the primary operating system. For synchronizing the timer of the primary operating system and as shown in fig. 4, besides the main timer (MT) in the system driver a virtual ti-
25 mer (VT), with a lower clock rate runs and which is implemented by the clock rate of the primary operating system (POS) as soon as the corresponding time has actually elapsed. If the secondary operating system does not deliver the computing time for a longer time period, the clock
30 for the primary operating system runs faster until the time difference has again been made up and specifically in the

timing of the SOS timer. Therefore the clock in the primary operating system does not run slow.

5

10

15

20

25

30